

# Real Time Big Data Architecture: Couchbase Server Certified for Hortonworks Data Platform

To harness the power of big data for real time business decision making, enterprises require a modern infrastructure that includes Hadoop and NoSQL. Couchbase Server is a high performance, distributed NoSQL database for enterprise web, mobile, and IoT applications. It's engineered to meet the high throughput, low latency requirements of enterprise operational workloads. Hortonworks Data Platform (HDP) is the Hadoop enterprise platform for batch, interactive, and real-time analytical workloads.

Couchbase Server is certified for Hortonworks Data Platform. By integrating Couchbase Server with HDP – utilizing the Kafka, Storm, and Sqoop components – enterprises can leverage offline and real-time stream based analytics to take action on data in motion.

## Kafka for Ingestion

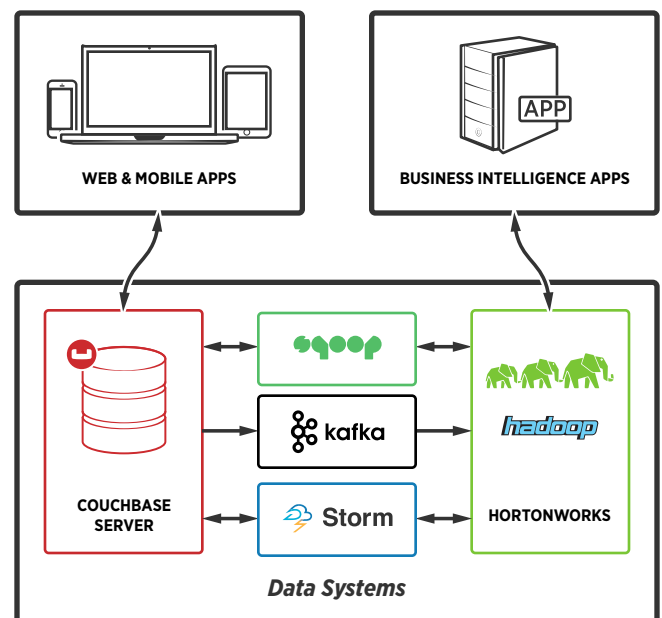
For use cases where continuous, one-way writes are required – such as telemetry readings from sensors – the Kafka producer can be utilized to publish writes to Couchbase Server in a message queue for high throughput, low latency ingestion. Kafka leverages the database change protocol, a unique architectural feature of Couchbase Server 3.0, to stream data from Couchbase Server to the message queue in real time. When data is inserted, updated, or deleted in Couchbase Server, it's automatically published to the message queue.

## Storm for Stream Processing

For use cases where stream processing is required – such as real-time financial market trend analysis – a Storm spout can be utilized to pull data into Storm from Couchbase Server for real-time analytics. When data is inserted, updated, or deleted in Couchbase Server, it's pulled into Storm. While a spout pulls data into Storm, a bolt processes or persists data. A Storm bolt can be utilized to persist data, raw or processed, to Couchbase Server for real-time access by reporting and visualization dashboards.

## Sqoop for Batch Import / Export

For use cases where batch import/export between Couchbase Server and Hadoop makes sense, the Sqoop plugin can be utilized to transfer data between Couchbase Server and HDP. It leverages map/reduce to import and export data via parallel tasks. It enables enterprises to export operational data from Couchbase Server to HDP for offline analytics and processing, and to import processed data from HDP into Couchbase Server to improve operational efficiency.



2440 West El Camino Real | Ste 600  
Mountain View, California 94040

1-650-417-7500

[www.couchbase.com](http://www.couchbase.com)

## About Couchbase

Couchbase delivers the world's highest performing NoSQL distributed database platform. Developers around the world use the Couchbase platform to build enterprise web, mobile, and IoT applications that support massive data volumes in real time. The Couchbase platform includes Couchbase Server, Couchbase Lite – the first mobile NoSQL database, and Couchbase Sync Gateway. Couchbase is designed for global deployments, with configurable cross data center replication to increase data locality and availability. All Couchbase products are open source projects. Couchbase customers include industry leaders like AOL, AT&T, Bally's, Beats Music, BSKyB, Cisco, Comcast, Concur, Disney, eBay, KDDI, Nordstorm, Neiman Marcus, Orbitz, PayPal, Rakuten / Viber, Tencent, Verizon, Wells Fargo, Willis Group, as well as hundreds of other household names. Couchbase investors include Accel Partners, Adams Street Partners, Ignition Partners, Mayfield Fund, North Bridge Venture Partners, and West Summit.